

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application of

Applicant : Steven K. Rogers

Title : USE OF COMPUTER-AIDED DETECTION SYSTEM OUTPUTS IN  
CLINICAL PRACTICE

Docket : QUA 001 IA

**MS PATENT APPLICATION**

Commissioner for Patents

PO Box 1450

Alexandria VA 22313-1450

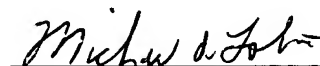
Sir:

EU 891072430 US

**"Express Mail" Mailing Label Number** EU891072430US

**Date of Deposit** September 29, 2003

**I hereby certify that this paper or fee is being deposited with the  
United States Postal Service "Express Mail Post Office to Addressee"  
service under 37 CFR 1.10 on the date indicated above and is  
addressed to MS PATENT APPLICATION, Commissioner for  
Patents, PO Box 1450, Alexandria VA 22313-1450**



Michael D. Folkerts

Reg. No. 33,348

Use of Computer-Aided Detection System Outputs in Clinical Practice

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation-in-part of U.S. Application No. 10/280,237, filed October 25, 2002, now allowed, which is a continuation of U.S. Application No. 09/938,908, filed August 24, 2001, now U.S. Patent No. 6,556,699, which is a continuation of U.S. Application No. 09/758,889, filed January 11, 2001, now U.S. Patent No. 6,389,157, which is a continuation of U.S. Application No. 09/416,437, filed October 12, 1999, now U.S. Patent No. 6,205,236, which is a continuation of U.S. Application No. 09/141,802, filed August 28, 1998, now U.S. Patent No. 5,999,639, which claims the benefit of U.S. Provisional Application No. 60/057,801, filed August 28, 1997, U.S. Provisional Application No. 60/066,996, filed November 28, 1997, and U.S. Provisional Application No. 60/076,760, filed March 3, 1998, the entire disclosures of all of which are incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to computer-aided detection (CAD) devices in medical imaging, with particular application to the use of CAD system outputs in subsequent procedures.

2. Discussion of Background

Mammography coupled with physical examination is the current standard for breast cancer screening. Although mammography is a well-studied and standardized methodology, for 10 to 30 percent of women diagnosed with breast cancer, their mammograms were interpreted as negative. Additionally, only 10 to 20 percent of patients referred for biopsy based on mammographic findings prove to have cancer. Further, estimates indicate the malignancies missed by radiologists are evident in two-thirds of the mammograms retrospectively. Missed detections may be attributed to several factors including poor image quality, improper patient positioning, inaccurate interpretation, fibroglandular tissue obscuration, subtle nature of radiographic findings, eye fatigue, or oversight.

Computer-aided detection systems are now widely used to assist in the early detection of disease. Three systems commercially available for the detection of signs of breast cancer in mammographic images are Second Look, CADx Systems, Beavercreek, OH;

ImageChecker, R2 Technologies, Sunnyvale, CA; MammoReader, iCAD, Hudson, NH. An additional system for detecting signs of lung cancer in chest images is RapidScan, Deus Technologies, Rockville, MD.

In practice, CAD is applied in the following fashion. A patient arrives at an imaging facility. A technologist uses imaging equipment to create medical images. These images are then input to a CAD system. The CAD system analyzes the images and provides a visual representation of the images, marking potentially suspicious regions on a paper or electronic display. Subsequently, when the radiologist is reading the images for the patient, the CAD outputs are examined to provide potentially new areas for consideration. The utility of CAD is in prompting a radiologist to consider areas not initially detected in a first reading of the images.

When a radiologist determines a suspicious region exists, further action is required. Common options include additional imaging procedures. In some cases, the suspicious area needs to be re-imaged under different positioning, compression, and magnification settings. A common element in these procedures is that a specific portion of the body needs to be identified and re-imaged. Currently, radiologists commonly indicate suspicious areas on film images by circling regions with a grease pencil. The marked films are then sent back to the technologist, who generates the requested images. Sometimes, multiple areas are marked on the film and the technologist must get clarification from the radiologist to ensure the proper area is re-imaged. In other cases, radiologists use a specialized paper form, Fig. 1, or electronic reporting system to indicate the location of the area requiring further study. An example in a mammographic application follows.

In a film-based clinic, the technologist receives a mammographic report describing the location of the area to be re-imaged and the original medical images marked with grease pencil. The current means used by radiologists to communicate with technologists have disadvantages. In mammographic reports, a radiologist describes the suspicious lesion characteristics and location to be re-imaged with words and sometimes an approximate sketch on the standardized line art breast outline, as shown in Fig. 2. An example description is “indeterminate microcalcifications in the upper outer quadrant of the right breast.” Although technologists can locate the specified area in the films most of the time, occasionally they will require further input from the radiologist or possibly even re-image the wrong area.

Disadvantages in the current reporting method are now described. When an approximate sketch is used to help specify the location of the suspicious lesion, a simple line art drawing of breast outlines may be provided on a special purpose form. The radiologist then indicates suspicious areas by marking the intended area on the outline. For paper and electronic reporting system, a simple and standard drawing of breast outlines is used for every patient. The radiologist marks locations of suspect regions to be re-imaged. The technologist interprets the radiologist's marks and refers to the original medical images to localize the proper area for the subsequent imaging procedure. When the particular patient's anatomy is not well represented by the standard line art drawing, the technologist must make a decision based personal experience to best localize the desired region.

One difficulty with marking films with a grease pencil is that in some instances, multiple areas of suspicion are marked without sufficient clarification regarding which location to re-image. The radiologist may create the additional marks during the initial reading session when a grease pencil is used to mark candidate locations during the visual search process. This problem results from using the search imagery as a reporting mechanism. That is, the radiologist searches the films and must consider many potentially suspicious areas. The grease pencil marks help the radiologist ensure that areas found throughout the visual search process are fully considered. When the technologist uses this marked imagery to guide a subsequent procedure, many of the marks made during the search are not relevant.

## SUMMARY OF THE INVENTION

The present invention provides for the use of CAD output displays for providing accurate representation of areas to undergo re-imaging procedures. Since the CAD output is not used during the initial reading, the radiologist does not mark it until a final determination is reached regarding subsequent procedures. Additionally, since the CAD output contains versions of the original imagery, the regions indicated by the radiologist are shown in the context of the particular anatomical detail for the given patient. This detail will assist the technologist in more efficiently and accurately locating the exact area for subsequent exams.

Therefore, it is an object of the current invention to provide a convenient method and system for a radiologist to indicate regions requiring further study directly on a CAD system output and a means for providing the CAD outputs to a technologist.

Other objects and advantages of the invention will be apparent from the following description, the accompanying drawings and the appended claims.

#### BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a simple line art drawing of a breast outline.

Fig. 2 shows an annotated breast outline drawing.

Fig. 3 is a flow diagram illustrating the automated system for the detection of clustered microcalcifications in a digital mammogram.

Figs. 4 and 5 are flow diagrams illustrating the autocropping method and system of the invention.

Figs. 6-12 are flow diagrams illustrating in more detail the autocropping method and system of the invention.

Fig. 13 is a flow diagram illustrating in greater detail the clustered microcalcification detector of the invention.

Fig. 14 is a schematic diagram illustrating a  $3 \times 3$  cross-shaped median filter of the invention.

Fig. 15 is a three-dimensional plot of a Difference of Gaussians (DoG) filter kernel.

Fig. 16 is a cross-sectional view through the center of the DoG filter kernel of Fig. 15.

Fig. 17 is a flow diagram illustrating the global thresholding portion of the microcalcification detection system.

Fig. 18 is a flow diagram illustrating the dual local thresholding of the invention.

Fig. 19 is a flow diagram illustrating combining the results of global and dual-local thresholding.

Fig. 20 is a flow diagram illustrating the sloping local thresholding of the invention.

Fig. 21 is a flow diagram illustrating the clustering method of the invention.

Fig. 22 is a schematic diagram illustrating the clustering method of the invention.

Fig. 23 is a flow diagram illustrating the feature computation process of the invention.

Fig. 24 is a flow diagram illustrating a classifier having one discriminant function per class.

Fig. 25 is a schematic diagram illustrating a multi-layer perceptron neural network for a two-class classifier.

Fig. 26 is a histogram of testing results after detection and classification.

Fig. 27 is a flow diagram illustrating the parameter optimization method of the invention.

Fig. 28 is a plot of a free response receiver operating characteristic curve of the invention before classifying detections.

Fig. 29 is a plot of a free response receiver operating characteristic curve of the invention after classifying detections.

Fig. 30 is a plot of probability density functions showing the relationship between the probabilities of false negative and false positive detections.

Fig. 31 is a plot of probability density functions showing the relationship between the probabilities of true negative and true positive detections.

Fig. 32 is a Venn diagram showing the relationship between radiologist and CAD system detections.

Fig. 33 is a flow diagram illustrating a method for incorporating computer-aided diagnosis detections with those of a human interpreter for optimal sensitivity.

Fig. 34 is an overview of a networked and CAD-enabled medical imaging facility.

Fig. 35 is an overview of the method for using CAD outputs to guide re-imaging procedures.

Fig. 36 shows CAD system outputs indicating CAD detected areas overlaid on images showing anatomical detail.

Fig. 37 shows radiologist produced marks in addition to the CAD detections drawn directly on the CAD output.

Fig. 38 shows an original CAD system output displayed on a monitor.

Fig. 39 shows the modified CAD system output displayed on a monitor, where the radiologist has marked the CAD system detections for subsequent imaging.

Fig. 40 shows a CAD system output including pulldown menus.

Fig. 41 shows an example of entries available from the pulldown menus.

Fig. 42 shows a CAD system output displayed on a monitor, where the CAD system provides three detections.

Fig. 43 shows a modified CAD system output displayed on a monitor, where the radiologist has removed a CAD detection from the RMLO image.

Fig. 44 shows another original CAD system output displayed on a monitor.

Fig. 45 shows a modified CAD system output displayed on a monitor, where the radiologist has added a CAD detection to the RCC image.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring now to the drawings, wherein like reference numerals designate identical or corresponding parts throughout the several views, and more particularly to Fig. 3 thereof, there is shown a flow diagram illustrating a sequence of steps performed in order to detect the locations of clusters of microcalcifications within a digital mammogram.

In a first step 100, a digital mammogram is obtained using hardware such as digital mammography systems, or by digitizing mammography films using laser or charge-coupled device (CCD) digitizers. In an optimized cropping step 200, a rectangular analysis region containing breast tissue is segmented from the digital mammogram image and a binary mask corresponding to the breast tissue is created for use in later processing steps to decrease the time required for processing the mammogram image. The binary mask is also used to limit detections to areas of the image containing breast tissue.

Clustered microcalcifications are detected in a clustered microcalcification detection step 300. After first filtering the cropped image with a median filter to reduce noise, the image is filtered using an optimized difference of Gaussians (DoG) filter to enhance the microcalcifications. The DoG-filtered image is then subjected to optimized threshold tests to detect potential microcalcifications. The detected microcalcifications are shrunk to single-pixel representations and detections outside of the breast area are removed. The remaining microcalcifications are grouped into clusters. Features are then computed for the clusters. Detected clusters are classified as either suspicious or non-suspicious in a classification step 400.

The parameters used by the autocropping, clustered microcalcification detection, and classification steps 200, 300, 400 are optimized in a parameter-optimizing step 500. The parameters are optimized by parameter-optimizing means that uses a genetic algorithm (GA) so

as to maximize the true-positive detection rate while minimizing the false-positive detection rate. Of course, other optimization schemes may be used as well.

The detected clustered microcalcifications are stored in a list of image coordinates. The detection results are processed in a processing step 600 by simply adding an offset to each of the microcalcification coordinates to account for translation of the coordinates incurred as a result of the cropping procedure. Detected clustered microcalcifications are indicated on the digital mammogram by means of rectangles drawn around the clustered microcalcifications in a display step 700. Other indicators may be used such as, for example, arrows pointing to suspected microcalcifications, or ellipses around suspected microcalcifications.

#### ACQUIRING A DIGITAL REPRESENTATION OF A MAMMOGRAM

One method of obtaining digital mammograms comprises digitizing radiological films by means of a laser or charge-coupled device (CCD) scanner. Digital images obtained in this manner typically have a sample spacing of about 100  $\mu\text{m}$  per pixel, with a gray-level resolution of 10 to 12 bits per pixel. In one embodiment of the present invention, radiological films are scanned using a Model CX812T digitizer manufactured by Radiographic Digital Imaging of Compton, California, to produce digital images having 50  $\mu\text{m}$  spacing per pixel and 12 bits of gray-level resolution per pixel. Another possible input source for digital images is a digital mammography unit from Trex Medical Corporation of Danbury, Connecticut, which has a spatial resolution of about 45  $\mu\text{m}$  per pixel and a gray-level resolution of 14 bits per pixel.

The digital images are stored as digital representations of the original mammogram images on computer-readable storage media. In a preferred embodiment, the digital representations or images are stored on a 2 GB hard drive of a general-purpose computer such as a PC having dual Pentium II  $\text{\textcircled{R}}$  microprocessors running at 200 MHZ, 512 MB of RAM memory, a ViewSonic PT813  $\text{\textcircled{R}}$  monitor, a pointing device, and a Lexmark Optra S1625  $\text{\textcircled{R}}$  printer. The system operates within a Windows NT  $\text{\textcircled{R}}$  operating system.

#### AUTOCROPPING

As may be seen in Figs. 4 and 5, a digital mammogram image 190 is first cropped to segment an analysis region 296 from the image and produce a binary mask 298 corresponding



to breast tissue in the analysis region. Preferably, the cropping is performed automatically, although it could be cropped manually. The image is cropped as a preliminary step because the breast tissue does not cover the whole radiographic film. Focusing the processing of the image on only that portion of the image which contains breast tissue reduces the time required to process the image. Also, other items appearing on the film, such as labels and patient information, are excluded from consideration, and false-positive indications lying outside of the breast tissue area are eliminated.

Referring to Figs. 6 through 12, the autocropping process will be described in detail. The image is first subsampled from 50  $\mu\text{m}$  to 400  $\mu\text{m}$  to reduce the amount of data to be processed in step 202. Of course, the image may be downsampled to other resolutions as desired. Not all of the original image data is needed to reliably segment the breast tissue from the remainder of the image. Subsampling every eighth pixel in both the horizontal and vertical directions reduces the amount of data by 64 times. For purposes of segmenting the breast tissue from the rest of the image, the consequent loss of resolution is immaterial.

A white border twenty pixels in width is added around all sides of the subsampled image in step 204. White corresponds to the maximum pixel value possible given the number of bits used to represent each pixel. For images having 12 bits of gray-scale resolution, the maximum gray-scale value is 4095. The bordered image is then thresholded in step 206 with a relatively high threshold value such that most of the breast tissue is guaranteed to be less than the threshold to produce a binary image. In one embodiment of the invention, the threshold is set equal to a predetermined percentage of the gray-scale value of a pixel near the top middle portion of the image. The thresholded image is then inverted, that is, ones become zeroes and zeroes become ones, in step 208. The inverted image is then dilated in step 210. Dilation is a morphological operation in which each pixel in a binary image is turned on, that is, set to a value of one, if any of its neighboring pixels are on. If the pixel is already on, it is left on.

In step 212 the dilated image is cropped to the size of the largest blob. Blobs are contiguous groups of pixels having the value one. This step 212 removes bright borders from the subsampled mammogram representation while ensuring that none of the breast area is reduced. Other techniques that threshold to find the border have a very difficult time dealing with bright areas in the breast adjacent to the border such as, for example, when breast implants are visible in the image. Pixels from the original image, resulting from step 202, corresponding to the

locations of the pixels in the cropped blob, are selected for subsequent processing. Note that this is a simple subset of pixels from the input image.

The image from step 212 is histogram equalized in step 214. The average brightness of the image will vary widely from mammogram to mammogram. Moreover, different digitizers having different optical density characteristics are an additional source of variability in brightness levels in the digital representation of the mammogram. The breast mask that is the output of the autocropper is mainly defined by means of a region-growing algorithm that requires a single contrast setting to work properly. However, it has been determined experimentally that a single contrast setting will not work for a wide range of image inputs. Therefore, each image is mapped into a normalized image space using an automatic histogram enhancement process, after which a single contrast setting works well.

First, a histogram of the image is obtained. Typically, most of the data in the breast area will be in the lower histogram bins (corresponding to gray-scale values of about 0 - 1000), with borders and labels being in the higher bins (corresponding to gray-scale values of about 4000 - 4095) for 12-bit data. The upper and lower bin values that contain the typical breast data are determined. The lower bin value is the first highest peak encountered when going from the lowest gray-scale value toward the highest gray-scale value. The upper bin is the last zero-value bin encountered when going from the highest gray-scale level toward the lowest gray-scale value. Then the data are reduced to an eight-bit representation and linearly stretched over the range of the data type. For example, values in the lower bins are set to zero. Values of data in the upper bins are set to 255. The rest of the data are then linearly mapped between the lower and upper bins.

After the image has been histogram equalized, the equalized image may be considered to be a matrix. The image matrix is divided into left and right halves, of equal size if possible, and the brighter side is selected in a step 216. The sums of all the pixels in the left and right halves are computed. The sum values are then compared and the side having the greater sum is the brighter side.

Prior to region growing the brighter side, algorithm variables are initialized in step 218. The size of the region-grown mask is preliminarily checked in step 220. If it is large enough, then the mask is acceptable. Otherwise, processing continues to find the mask. The side of the image to be region grown is selected in step 222. In step 224 this region is searched to

find its maximum gray-scale value. This maximum value is used to find a pixel to start a region-growing algorithm. Region growing is the process of grouping connected pixels sharing some like characteristic. The choice of characteristic influences the resultant region. The input to a region growing function is a gray-scale image and a starting point to begin growing. The output is a binary image with ones indicating pixels within the grown region, *i.e.*, blobs. Region growing will create a single blob, but that blob may have within it internal holes, that is, pixels that are off. To grow a blob, each of the four nearest neighbors of a pixel of interest are looked at. The contrast ratio is computed for each nearest neighbor pixel. If the contrast ratio is less than a contrast ratio threshold, then the neighbor pixel is set to a one in a binary mask image. Otherwise, the neighbor pixel is set to zero. The region growing algorithm spirals outwardly from the starting or seed pixel, progressively looking at nearest neighbor pixels until done. To those skilled in the art, it is clear that other region growing algorithms may also be applied.

In step 226, region growing begins with the pixel identified from the previous step 224 to produce a binary mask. The size of the mask resulting from step 226 is computed in step 228 and checked in step 230. There may be three points of failure for this approach. First, the brightest point in the search region may be an artifact outside the breast. Therefore, if the resulting mask is not large enough (50 pixels), then the search region is moved closer to the side of the image and searched again. This is repeated three times, each time lowering the contrast value threshold. This corresponds to the path taken through steps 232 and 234. Second, the side selection approach may be in error. Therefore, if a valid breast mask is not found in the first side searched, then the other side of the equalized image is searched. This corresponds to the path taken through steps 236 and 238. Third, if a valid breast mask is not found on either side, then the whole breast is thresholded and the largest object is taken to be the breast mask in step 240.

Since a constant contrast value is used in the region-growing algorithm, some masks will be too large. Typically, there will be “tails” along the edge of the digitized mammogram image where extra light leaked in while the original mammogram film was being digitized. The tails are reduced by applying a series of erodes and then a series of dilates to the image. Erosion is a morphological operation in which each pixel in a binary image is turned off unless all of its neighbors are on. If the pixel is already off, it is left off. But first, the holes in the mask must be filled in or the multiple erodes may break the mask into disjoint sections. Thus, holes in the mask are closed in step 242 by means of a majority operation. The majority

operation is a morphological operation in which each pixel in a binary image is turned on if a majority of its neighboring pixels are on. If the pixel is already on, it is left on.

However, another problem is that some smaller breast masks can not undergo as many erodes as can larger breast masks. Therefore, as a fail-safe measure, the sum of the breast mask is taken before and after the erodes and dilates. If the size is reduced too much (*i.e.*, by more than 50%), the original mask before the morphological operators is used. Thus, a duplicate copy of the mask is made in step 244 before the mask is eroded and dilated in steps 246 and 248, respectively. The size of the resultant mask is then computed in step 250 and compared with the size of the mask from step 242 in step 252. If the new size is less than half the old size, then the duplicate mask, from step 244, is selected in step 254 for subsequent processing. Otherwise, the resultant mask from step 248 is used.

The original image (from step 202) is then cropped to the size of the breast mask just found (either from step 242 or step 248) in step 256. In case the resulting mask is too small for subsequent processing, a crop adjustment is always made in step 258. The adjustment comes in the form of increasing the size of the breast mask bounding box by including additional pixels from the original image in the cropped image.

The cropped image is then automatically histogram enhanced in step 260 as previously described above in connection with step 214. This enhanced image is passed through a loose region growing step 262 to produce a generous mask. This means that the image is subjected to a lower threshold to yield more “on” pixels. This mask is then subjected to hole-closing, eroding, and dilating in steps 264, 266, and 268, respectively, as above, but to a lesser degree.

The same steps described above are repeated one final time in steps 270 through 276, but the crop adjustments are less and the contrast value is increased for a tight region growing step 276. This tight region growing step 276 can afford the higher contrast value since it will be region growing in just the cropped image. This results in a parsimonious estimate of breast tissue. The resulting mask is segmented to find the largest object in step 278 and its bounding box shrunk to just enclose the object in step 280. There may still be some holes in the breast mask. Therefore, after crop adjustments in step 282, the mask is inverted in step 284 and the largest object is found in step 286. This largest object is extracted and then inverted in step 288 to obtain the penultimate mask.

The final mask is obtained by closing holes in the penultimate mask with multiple majority operations and dilations in step 290. The image is then cropped to the size of the resulting mask and the autocropping is complete. An important result from the autocropper is the offset of the cropped image. This is the pixel location in the original image that corresponds to the pixel in the upper left pixel of the cropped image. Keeping track of all the cropping and crop adjustments determines this offset value.

The output of the autocropping process is a rectangular array of pixels representing a binary mask wherein the pixels corresponding to breast tissue are assigned a value of one while the remainder of the pixels are assigned a value of zero. Put another way, the binary mask is a silhouette of the breast made up of ones while the background is made up of zeroes.

Parameters of the autocropper may be optimized to obtain better breast masks. The procedure is described below in the optimization section.

## DETECTION OF CLUSTERED MICROCALCIFICATIONS

Turning now to Fig. 13, there is seen therein a flow diagram illustrating in greater detail the clustered microcalcification detection system 300 of the invention.

That portion of the digital representation of the mammogram corresponding to the analysis region 296, designated a cropped sub-image 302, produced in the cropping step 200, is first processed to reduce noise in a noise reduction step 310 to reduce digitization noise that contributes to false detections of microcalcifications. The noise-reduced image is then filtered using an optimized target-size-dependent difference of Gaussians (DoG) spatial kernel in step 320 to enhance differences between targets and background, thus creating global and local maxima in the filtered image. The optimized DoG-filtered image is then thresholded in step 340 to segment maxima that represent potential detections of microcalcifications.

The detected maxima are converted to single-pixel coordinate representations in a conversion step 350. The coordinate representations of the detected maxima are compared with the binary mask of the analysis area in a first false-positive removal step 360 to remove false detections outside the breast mask area. The remaining coordinate representations in the analysis area are clustered in a clustering step 370. Features are computed for the remaining clusters in a feature computation step 380 and used to remove non-suspicious detections in a classifying step

400 (Fig. 3). The remaining detections are outputted as detected clustered microcalcifications in an outputting step 600 in the form of cluster coordinates.

Turning now to a more detailed discussion of the steps in the clustered microcalcification detection process, the digital mammogram image is first filtered to reduce noise in the image. Although the main limitation in image quality should be the granularity of the film emulsion, noise is introduced from the process of digitization. This noise may later be detected as a pseudocalcification. In this system, a cross-shaped median filter is used because it is well known to be extremely effective at removing single-pixel noise. The median filter is a non-linear spatial filter that replaces each pixel value with the median of the pixel values within a kernel of chosen size and shape centered at a pixel of interest. Referring to Fig. 14, it may be seen that the cross shape is formed by the set of pixels which include the center pixel and its four nearest neighbors. The cross shape preserves lines and corners better than typical block-shaped median filters and limits the possible substitution to the four nearest neighbors, thereby reducing the potential for edge displacement.

After noise has been reduced, the image is filtered with an optimized DoG kernel to enhance microcalcifications. Filtering is accomplished by convolving the noise-reduced image with the DoG kernel. In an alternative embodiment, filtering is accomplished by first obtaining the fast Fourier transforms (FFTs) of the noise-reduced image and the DoG kernel, then multiplying the FFTs together, and taking the inverse FFT of the result.

The DoG kernel was chosen because neurophysiological experiments provide evidence that the human visual pathway includes a set of “channels” that are spatial frequency selective. Essentially, at each point in the visual field, there are size-tuned filters or masks analyzing an image. The operation of these spatial receptive fields can be approximated closely by a DoG.

The 2-D Gaussian mask is given as:

$$G(x, y) = ce^{\frac{-(x^2+y^2)}{2\sigma^2}}$$

where  $c$  normalizes the sum of mask elements to unity,  $x$  and  $y$  are horizontal and vertical indices, and  $\sigma$  is the standard deviation. Using Equation 1, the difference of two Gaussians with different  $\sigma$  yields:

$$DoG(x, y) = c_1 e^{\frac{-(x^2+y^2)}{2\sigma_1^2}} - c_2 e^{\frac{-(x^2+y^2)}{2\sigma_2^2}}$$

It has been shown that when  $\sigma_2 = 1.6\sigma_1$ , then the DoG filter's response closely matches the response of human spatial receptive filters. Therefore, with motivation from human physiology, let the ratio of the DoG standard deviation constants be 1:1.6. Then, for a target of size (average width)  $t$  pixels, use  $\sigma_2 = t/2$  and, from the rule of thumb,  $\sigma_1 = \sigma_2/1.6$ .

Since microcalcifications typically range from 100 to 300  $\mu\text{m}$  in diameter, potential target sizes for the 50  $\mu\text{m}$  digitized mammograms correspond to 2 to 6 pixels. It has been found that a DoG kernel constructed using an optimization technique for selecting the target size parameter, such as the GA detailed below, has an optimized target size of  $t = 6.01$  pixels. The targetsize  $t$  will vary depending on such factors as the resolution and scale of the image to be processed. The impulse response of a DoG filter having  $t = 6.01$  pixels and  $\sigma_2 = 1.6\sigma_1$  is shown in Figs. 15 and 16.

Once the noised-reduced cropped image has been DoG filtered to enhance differences between targets and background, the DoG-filtered subimage contains differences in gray levels between potential microcalcifications and background. Although microcalcifications tend to be among the brightest objects in DoG-filtered subimages, they may exist within regions of high average gray levels and thus prove difficult to reliably segment. The thresholding process used in one embodiment of the invention that generally addresses these concerns involves pair-wise pixel "ANDing" of the results of global histogram and locally adaptive thresholding. However, the preferred embodiment of the invention uses sloping local thresholding.

Since targets tend to exist within an image's higher gray levels, then the global threshold may be approximated by finding the level which segments a preselected percentage of the corresponding higher pixel levels in the image histogram. An embodiment of a global thresholding method is illustrated in Fig. 17. Locally adaptive thresholding may be implemented by varying the high and low thresholds based on the local pixel value mean and standard deviation. An embodiment of a dual-local thresholding method is illustrated in Fig. 18.

After computing the image histogram,  $p(r_k)$ , the gray level threshold,  $g$ , used to segment a preselected upper fraction,  $f$ , of the histogram, is found using:

$$f = 1 - \sum_{k=0}^g p(r_k)$$

where  $r_k$  is the  $k^{\text{th}}$  gray level,  $0 \leq g \leq g_{\max}$ , and  $g_{\max}$  is the maximum gray level in the image.

The locally adaptive thresholds,  $t_{lo}$  and  $t_{hi}$ , are found using

$$t_{lo} = k_{lo} \sigma_{NN}(x, y) + \mu_{NN}(x, y)$$

and

$$t_{hi} = k_{hi} \sigma_{NN}(x, y) + \mu_{NN}(x, y)$$

where  $k_{lo}$  and  $k_{hi}$  are used to preselect the multiple of  $\sigma_{NN}(x, y)$ , the local standard deviation of gray-level intensities, and  $\mu_{NN}(x, y)$  is the local gray-level mean of the  $N \times N$  neighborhood centered on the pixel at  $(x, y)$  of the DoG-filtered image. Other neighborhood shapes, such as rectangular, circular, and ellipsoidal, may also be used. Pixels whose brightness or gray-level value falls within the threshold interval, that is,  $t_{lo} < \text{brightness} < t_{hi}$ , are set equal to one.

Optimization of  $f$ ,  $k_{lo}$ ,  $k_{hi}$ , and  $N$  is discussed below in connection with the parameter-optimizing process. The results of the global thresholding process may be combined with the results of the local thresholding step by logically ANDing them as shown in Fig. 19. Alternatively, either thresholding method may be used alone.

The preferred thresholding means are illustrated in Fig. 20 wherein it may be seen that an  $N \times N$  window is centered at a pixel  $x, y$  in the input image  $p(x, y)$ . The mean,  $\mu(x, y)$ , and standard deviation,  $\sigma(x, y)$ , of the digital mammogram image pixels under the window are computed. A local threshold value,  $T(x, y)$ , is computed as:

$$T(x, y) = A + B \mu(x, y) + C \sigma(x, y)$$

where values for  $N$ ,  $A$ ,  $B$ , and  $C$  are computed during a parameter optimization stage, discussed below. Values for  $T(x, y)$  are computed for every  $x, y$  location in the image.

The digital mammogram has also been DoG filtered, producing an image  $d(x, y)$ . Each pixel of the DoG-filtered image  $d(x, y)$  is compared to the threshold value  $T(x, y)$ . Pixels in the locally thresholded image  $l_s(x, y)$  are set to one where values of the DoG-filtered image are greater than the threshold, and set to zero elsewhere.

The advantage of this novel local sloping thresholding method over prior art thresholding methods is that the threshold is computed from the pixels in a pre-DoG-filtered



image rather than from a post-DoG-filtered image. This eliminates the need for background trend correction. In conventional local thresholding, the threshold is computed as:

$$T(x, y) = B \mu(x, y) + C \sigma(x, y)$$

from the mean and standard deviation of the DoG-filtered image. The problem of using a local threshold computed from the DoG-filtered image is that DoG-filtered images typically have mean values close to zero and standard deviations significantly affected by the presence of targets.

Local thresholds computed from the statistics of the DoG-filtered image suffer from the following adverse effects. First, since the mean value is close to zero, a degree of freedom is lost in the computation of the threshold, which becomes essentially a function of the standard deviation. Second, the absolute brightness of the input image is lost. To keep many spurious detections from occurring, it is desirable to have high thresholds in bright regions. However, the information about the local mean of the input image is not available in the DoG-filtered image. Finally, the standard deviations of DoG-filtered images are increased by detections of targets. This is so because when local bright spots of proper size exist in the original image, large gray-scale values result in the DoG-filtered image. Thus, the presence of targets in a region increases the local standard deviation thereby raising the threshold of that region. The higher threshold reduces the probability of passing a bright spot to subsequent processing stages.

The novel local thresholding method just described solves the above problems by computing thresholds from the input image, which are then applied to the DoG-filtered image. Additionally, the threshold computed here includes an offset term  $A$ , which is independent of the local image mean.

After thresholding, detections are converted to single-pixel representations by computing the centroid or center of gravity of groups of contiguous pixels found by the thresholding process. Detections are thus represented as single pixels having a value of logical one while the remaining pixels have a value of logical zero.

False-positive detections outside of the breast area are removed by logically ANDing the binary mask from the autocropper with the single-pixel representations of the detections.

Calcifications associated with malignancies usually occur in clusters and can be extensive. The cluster detection module identifies clusters based on a clustering algorithm as depicted in Fig. 21. Specifically, a suspicious cluster is declared when at least  $\mu Cs_{min}$  or more detected signals are separated by less than a nearest neighbor distance,  $d_{nn}$ . Optimization of  $\mu Cs_{min}$  and  $d_{nn}$  is discussed below in connection with the parameter optimizing process. Fig. 22 illustrates the clustering process for the case wherein  $\mu Cs_{min} = 5$  and  $d_{nn} = 4$ .

Additional false-positive clustered microcalcifications are removed by means of a classifier, detailed below. Features are extracted for each of the potential clustered microcalcifications as shown in Fig. 23. The eight features computed for each of the potential clustered microcalcifications in a preferred embodiment are:

1. The larger eigenvalue ( $\lambda_1$ ) of the covariance matrix of the points in a cluster;
2. The smaller eigenvalue ( $\lambda_2$ ) of the covariance matrix of the points in a cluster;
3. The ratio of the smaller eigenvalue of the covariance matrix to the larger eigenvalue of the covariance matrix of the points in a cluster. Equivalent to the ratio of the minor axis to the major axis of an ellipse fitted to cover the points in a cluster;
4. Linear density calculated as the number of detected microcalcifications divided by the maximum interpoint distance;
5. Standard deviation of the distances between points in a cluster;
6. Mean minus median of the distances between points in a cluster;
7. Range of points in cluster calculated as maximum interpoint distance minus the minimum interpoint distance; and
8. Density of a cluster calculated as the number of detections divided by the area of a box just large enough to enclose the detections.

Of course, other features could be computed for the potential microcalcification clusters, and the invention is not limited to the number or types of features enumerated herein.

## CLASSIFYING DETECTIONS

The cluster features are provided as inputs to the classifier, which classifies each potential clustered microcalcification as either suspicious or not suspicious. In practice, the clustered microcalcification detector is only able to locate regions of interest in the digital

representation of the original mammogram that may be associated with cancer. In any detector, there is a tradeoff between locating as many potentially suspicious regions as possible versus reducing the number of normal regions falsely detected as being potentially suspicious. CAD systems are designed to provide the largest detection rates possible at the expense of detecting potentially significant numbers of regions that are actually normal. Many of these unwanted detections are removed from consideration by applying pattern recognition techniques.

Pattern recognition is the process of making decisions based on measurements. In this system, regions of interest or detections are located by a detector, and then accepted or rejected for display. The first step in the process is to characterize the detected regions. Toward this end, multiple measurements are computed from each of the detected regions. Each measurement is referred to as a feature. A collection of measurements for a detected region is referred to as a feature vector, wherein each element of the vector represents a feature value. The feature vector is input to a discriminant function.

Referring to Fig. 24, there may be seen therein a classifier having a feature vector  $x$  applied to a set of discriminant functions  $g(x)$ . The classifier shown in Fig. 24 is designed with one discriminant function per class. A discriminant function computes a single value as a function of an input feature vector. Discriminant functions may be learned from training data and implemented in a variety of functional forms. The output of a discriminant function is referred to as a test statistic. Classification is selecting a class according to the discriminant function with the greatest output value. The test statistic is compared to a threshold value. For values of the test statistic above the threshold, the region or detection associated with the feature vector is retained and displayed as potentially suspicious. When the test statistic is below the threshold, the region is not displayed.

Many methods are available for designing discriminant functions. One approach considered for this invention is a class of artificial neural networks. Artificial neural networks require training, whereby the discriminant function is formed with the assistance of labeled training data.

In a preferred embodiment, the classification process is implemented by means of a multi-layer perceptron (MLP) neural network (NN). Of course, other classifier means could be used such as, for example, a statistical quadratic classifier. Only potential clustered microcalcifications classified as suspicious are retained for eventual designation for a radiologist.

Alternatively, it may be desirable to iteratively loop between MLP NN analysis of the individual microcalcification detections and the microcalcification clusters.

Referring to Fig. 25, a schematic diagram of an MLP NN may be seen therein. The MLP NN includes a first layer of  $J$  hidden layer nodes or perceptrons 410, and one output node or perceptron 420 for each class. The preferred embodiment of the invention uses two output nodes, one each for the class of suspicious detections and the class of non-suspicious detections. Of course, more or fewer classes could be used for classifying clusters of microcalcifications. Each computed feature  $x_i$  is first multiplied by a weight  $w_{i,j}$ , where  $i$  is an index representing the  $i^{\text{th}}$  feature vector element, and  $j$  is an index representing the  $j^{\text{th}}$  first layer node. The output  $y_j$  of each first layer perceptron 410 is a nonlinear function of the weighted inputs and is given by:

$$y_j = f \left( \sum_{i=1}^d (w_{i,j} \times x_i) \right)$$

where  $d$  represents the total number of features  $x_i$  and  $f(\bullet)$  is typically a saturating nonlinearity.

In this embodiment,  $f(\bullet) = \tanh(\bullet)$ . The first layer or hidden layer node outputs  $y_j$  are then multiplied by a second layer of weights  $u_{j,k}$  and applied to the output layer nodes 420. The

$$z_k(y) = f \left( \sum_{j=1}^J (u_{j,k} \times y_j) \right)$$

output of an output layer node 420 is a nonlinear function of the weighted inputs and is given by: where  $k$  is an index representing the  $k^{\text{th}}$  output node.

The hyperbolic tangent function is used in a preferred embodiment of the system because it allows the MLP NN to be trained relatively faster as compared to other functions. However, functions other than the hyperbolic tangent may be used to provide the outputs from the perceptrons. For example, linear functions may be used, as well as smoothly varying nonlinear functions, such as the sigmoid function.

The weight values are obtained by training the network. Training consists of repeatedly presenting feature vectors of known class membership as inputs to the network. Weight values are adjusted with a back propagation algorithm to reduce the mean squared error

between actual and desired network outputs. Desired outputs of  $z_1$  and  $z_2$  for a suspicious input are +1 and -1, respectively. Desired outputs of  $z_1$  and  $z_2$  for non-suspicious inputs are -1 and +1, respectively. Other error metrics and output values may also be used.

In this embodiment of the system, the MLP NN is implemented by means of software running on a general-purpose computer. Alternatively, the MLP NN could also be implemented in a hardware configuration by means readily apparent to those with ordinary skill in the art.

After training, each detected clustered microcalcification is classified as either suspicious or not suspicious by means forming the difference  $z_1 - z_2$ , then comparing the difference to a threshold,  $\theta$ . For values of  $z_1 - z_2$  greater than or equal to the threshold  $\theta$ , i.e.,  $z_1 - z_2 \geq \theta$ , the classifier returns a value of +1 for suspicious clustered microcalcifications, and for values of  $z_1 - z_2 < \theta$ , the classifier returns a value of -1 for non-suspicious clustered microcalcifications.

In order to arrive at optimum values for the respective weights, and the number of first layer nodes, the MLP NN was trained with a training set of feature vectors derived from a database of 978 mammogram images.

To develop and test the CAD system of the invention, truth data was first generated. Truth data provides a categorization of the tissue in the digital images as a function of position. Truth data was generated by certified radiologists marking truth boxes over image regions associated with cancer. In addition to the mammogram images, the radiologists also had access to patient histories and pathology reports.

The radiologists identified 57 regions of interest, containing biopsy-confirmed cancers associated with clustered microcalcifications, by means of truth boxes. All 978 images were then processed by the microcalcification detector of the invention to produce a plurality of feature vectors, a subset of which were associated with the 57 truth boxes. Half of the subset feature vectors were randomly chosen, along with about three times as many feature vectors not associated with clustered microcalcifications, to comprise the training set of feature vectors. The MLP NN, having a predetermined number of hidden nodes, was then trained using the training set. The remaining feature vectors were used as a test database to evaluate the performance of the MLP NN after training. Training of the MLP NN was carried out by means of the

Levenberg-Marquardt back propagation algorithm.

Alternatively, the MLP NN can be trained with other learning algorithms and may have nonlinearities other than the hyperbolic tangent in either or both layers. In an alternative embodiment with sigmoidal output nodes, the Bayes optimal solution of the problem of classifying clustered microcalcification detections as either suspicious or non-suspicious may be obtained.

In one run of the preferred embodiment during testing, before application of the MLP NN classifier to eliminate false-positive clustered microcalcifications, the detection procedure found about 93% of the true-positive clustered microcalcifications in both the training and test databases while indicating about 10 false-positive clustered microcalcifications per image. It was found that after an MLP NN classifier having 25 first layer nodes was used with the respective optimum weights found during training, 93% of the true-positive detections were retained while 57% of the false-positive detections were successfully removed. Referring to Fig. 26, there may be seen a histogram of the results of testing on the testing database after classification by the MLP NN. Of course, the MLP NN of the invention may be operated with more or fewer first layer nodes as desired.

### DISPLAYING DETECTIONS

After the locations of clustered microcalcifications have been determined, they are indicated on the original digitized mammogram image, or a copy of the original image, by drawing rectangular boxes around microcalcifications. Other means for indicating the locations of microcalcifications may be used, such as, for example, placing arrows in the image pointing at detections or drawing ellipses around the detections.

The locations of clustered microcalcifications are passed to the display detections procedure as a list of row and column coordinates of the upper left and lower right pixels bounding each of the clusters. The minimum row and column coordinates and maximum row and column coordinates are computed for each cluster. Bounding boxes defined by the minimum and maximum row and column coordinates are added to the original digitized image, by means well known in the art. The resulting image is then stored as a computer-readable file, displayed on a monitor, or printed as a hard-copy image, as desired.

In one embodiment of the system, the resulting image is saved to a hard disk on a

general-purpose computer having dual Pentium II ® processors and running a Windows NT ® operating system. The resulting image may be viewed on a VGA or SVGA monitor, such as a ViewSonic PT813 ® monitor, or printed as a hard-copy gray-scale image using a laser printer, such as a Lexmark Optra S1625 ®. Of course, other hardware elements may be used by those with ordinary skill in the art.

### OPTIMIZING THE PARAMETERS

Genetic algorithms (GAs) have been successfully applied to many diverse and difficult optimization problems. A preferred embodiment of this invention uses an implementation of a GA developed by Houck, et al. ("A Genetic Algorithm for Function Optimization," Tech. Rep., NCSU-IE TR 95-09, 1995), the entire disclosure of which is incorporated by reference herein, to find promising parameter settings. The parameter optimization process of the invention is shown in Fig. 27. This is a novel application of optimization techniques as compared to current computer-aided diagnosis systems that require hand tuning by experiment.

GAs search the solution space to maximize a fitness (objective) function by use of simulated evolutionary operators such as mutation and sexual recombination. In this embodiment, the fitness function to be maximized reflects the goals of maximizing the number of true-positive detections while minimizing the number of false-positive detections. GA use requires determination of several issues: objective function design, parameter set representation, population initialization, choice of selection function, choice of genetic operators (reproduction mechanisms) for simulated evolution, and identification of termination criteria.

The design of the objective function is a key factor in the performance of any optimization algorithm. The function optimization problem for detecting clustered microcalcifications may be described as follows: given some finite domain,  $D$ , a particular set of cluster detection parameters,  $x = \{t, f, k_{lo}, k_{hi}, N, \mu C s_{min}, d_{nn}\}$  where  $x \in D$ , and an objective function  $f_{obj} : D \rightarrow \mathbb{R}$ , where  $\mathbb{R}$  denotes the set of real numbers, find the  $x$  in  $D$  that maximizes or minimizes  $f_{obj}$ . When sloping local thresholding is used in the cluster detector, the parameters  $N$ ,  $A$ ,  $B$ , and  $C$  are optimized. Radiologic imaging systems may be optimized to maximize the TP rate subject to the constraint of minimizing the FP rate. This objective may be recast into the functional form shown in the following equation:

$$f_{obj}(x) = \begin{cases} -FP(x), & TP(x) \geq TP_{min} \\ FP_{penalty}, & \text{otherwise} \end{cases}$$

where maximization is the goal. For a particular set of cluster detection parameters, if the minimum acceptable TP rate,  $TP_{min}$ , is exceeded, the objective function returns the negative of the FP rate. Otherwise, if the TP rate falls below  $TP_{min}$ , the objective function returns a constant value,  $FP_{penalty} = -10$ . Other objective functions may also be used.

Since a real-valued GA is an order of magnitude more efficient in CPU time than the binary GA, and provides higher precision with more consistent results across replications, this embodiment of the invention uses a floating-point representation of the GA.

This embodiment also seeds the initial population with some members known beforehand to be in an interesting part of the search space so as to iteratively improve existing solutions. Also, the number of members is limited to twenty so as to reduce the computational cost of evaluating objective functions.

In one embodiment of the invention, normalized geometric ranking is used, as discussed in greater detail in Houck, et al., *supra*, for the probabilistic selection process used to identify candidates for reproduction. Ranking is less prone to premature convergence caused by individuals that are far above average. The basic idea of ranking is to select solutions for the mating pool based on the relative fitness between solutions. This embodiment also uses the default genetic operation schemes of arithmetic crossover and nonuniform mutation included in Houck, et al.'s GA.

This embodiment continues to search for solutions until the objective function converges. Alternatively, the search could be terminated after a predetermined number of generations. Although termination due to loss of population diversity and/or lack of improvement is efficient when crossover is the primary source of variation in a population, homogeneous populations can be succeeded with better (higher) fitness when using mutation. Crossover refers to generating new members of a population by combining elements from several of the most fit members. This corresponds to keeping solutions in the best part of the search space. Mutation refers to randomly altering elements from the most fit members. This allows the algorithm to exit an area of the search space that may be just a local maximum. Since restarting populations that may have converged proves useful, several iterations of the GA are



run until a consistent lack of increase in average fitness is recognized.

Once potentially optimum solutions are found by using the GA, the most fit GA solution may be further optimized by local searches. An alternative embodiment of the invention uses the simplex method to further refine the optimized GA solution.

The autocropping system may also benefit from optimization of its parameters including contrast value, number of erodes, and number of dilates. The method for optimizing the autocropper includes the steps of generating breast masks by hand for some training data, selecting an initial population, and producing breast masks for training data. The method further includes the steps of measuring the percent of overlap of the hand-generated and automatically-generated masks as well as the fraction of autocropped breast tissue outside the hand-generated masks. The method further comprises selecting winning members, generating new members, and iterating in a like manner as described above until a predetermined objective function converges.

In Figures 28 and 29, there may be seen therein free response receiver operating characteristic curves for the system of the invention for the outputs of the optimized microcalcification detector and the classifier, respectively. Fig. 28 represents the performance of the optimized detector before classifying detections, while Fig. 29 represents the performance of the system after classifying detections.

Although the GA has been described above in connection with the parameter optimization portion of the preferred embodiment, other optimization techniques are suitable such as, for example, response surface methodology. Of course, processing systems other than those described herein may be optimized by the methods disclosed herein, including the GA.

#### INCORPORATING CAD SYSTEM OUTPUTS FOR OPTIMAL SENSITIVITY

Performance metrics for detection of suspicious regions associated with cancer are often reported in terms of sensitivity and specificity. Sensitivity measures how well a system finds suspicious regions and is defined as the percentage of suspicious regions detected from the total number of suspicious regions in the cases reviewed. Sensitivity is defined as:

$$\text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}}$$

where TP is the number of regions reported as suspicious by a CAD system that are associated with cancers, and FN is the number of regions that are known to be cancerous that are not reported as suspicious. Specificity measures how well the system reports normal regions as normal. Specificity is defined as:

$$\text{Specificity} = \frac{\text{TN}}{\text{FP} + \text{TN}}$$

where TN represents regions correctly identified as not suspicious and FP represents regions reported as suspicious that are not cancerous.

Current CAD systems increase specificity by reducing FP. However, FP and TP are coupled quantities. That is, a reduction of FP leads to a reduction of TP. This implies that some of the suspicious regions that could have been detected are missed when the objective is to maintain high specificity.

Figures 30 and 31 illustrate relationships between the quantities TP, FP, TN, and FN. A measurement from a screening mammography image is represented by test statistic,  $x$ . The probability density function of  $x$  is represented by  $p(x)$  and the decision threshold is represented by  $\theta$ . If  $x$  is greater than  $\theta$ , a suspicious region is reported. Areas under the probability density functions represent probabilities of events. From Fig. 30 observe that increasing the threshold reduces the probability of FP decisions. However, observe from Fig. 31 that increasing the threshold simultaneously reduces the probability of TP decisions.

Another metric that exists for CAD systems is positive predictive value (PPV), which is defined as the probability that cancer actually exists when a region of interest is labeled as suspicious. PPV can be calculated from the following equation:

$$\text{PPV} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

Note that increasing TP or reducing FP increases PPV.

Radiologists and computers find different suspicious regions. Fig. 32 is a Venn diagram depicting a possible distribution of suspicious regions for man and machine detections. Some suspicious regions are found solely by a human interpreter or radiologist, some solely by a CAD system, some are found by both, and some are not found by either.

Referring to Fig. 33, there may be seen a preferred method for incorporating the outputs of a CAD system, and more particularly for the CAD system of the invention, with the observations of a human interpreter of a screening mammography image 10 for optimal sensitivity, wherein a radiologist examines the screening mammography image 10 in a step 20 and reports a set of suspicious regions 30 designated as S1. The CAD system then operates on the image 10 in a step 40 and reports a set of suspicious regions 50 designated as S2. The radiologist then examines set S2 and accepts or rejects members of set S2 as suspicious in a step 60, thereby forming a third set of suspicious regions 70 denoted as S3, which is a subset of S2. The radiologist then creates in a step 80 a set of workup regions 90 denoted as S4 which is the union of sets S1 and S3. The workup regions 90 are then recommended for further examination such as taking additional mammograms with greater resolution, examining the areas of the breast tissue corresponding to the workup regions by means of ultrasound, or performing biopsies of the breast tissue.

The present invention provides a method and system for using CAD system outputs to guide subsequent examinations.

#### USE OF CAD SYSTEM OUTPUTS TO GUIDE SUBSEQUENT PROCEDURES

Fig. 34 shows an overview of a networked and CAD enabled medical imaging facility. A technologist operates imaging equipment, 1100, to create medical images of a patient. The equipment may produce film or direct digital images. In case of film, a film digitizer, 1200, creates digital images. The digital imagery is then transferred across a network, 1300, to an Image and CAD output database, 1400. The technologist controls subsequent operations with a workstation, 1500. Simultaneous with or subsequent to the storage of the medical imagery in the database, computer-aided detection algorithms are applied using a CAD processor, 1600. The CAD system produces detections that may be stored with the corresponding imagery in the database. A printer, 1700, may alternatively print the CAD detections. A radiologist views the original medical imagery on image viewing equipment, such as a lightbox or radiologist

workstation, 1800, to begin reading the images. In an alternative embodiment, the CAD processing is accomplished on the radiologist workstation. Network connections to the radiologist and technologists are accomplished with wired or wireless means. When the radiologist deems a certain area suspicious, a common next step is to obtain additional images of that area. This invention relates to an improved method for communicating the location of the particular area to be re-imaged between the radiologist and the technologist.

Referring to Fig. 35, the sequence of steps in a clinical application of the invention is now described. A patient arrives at a facility where a technologist creates medical imagery, 1150. Types of imagery include mammography and chest x-ray films or full-field digital images. Other imaging procedures and anatomical regions may be imaged and still be within the scope of this invention. The images in the study commonly include approximately orthogonal views. In mammography for example, a craniocaudal (CC) and mediolateral oblique (MLO) view is created for each breast while in chest x-ray, lateral and posteroanterior views are created.

The imagery is annotated with patient identification information and possibly other associated process details such as date, time, technologist name, and so on. In a film-based clinic, a technologist or other specialist loads the films into a digitizer. The resultant imagery is analyzed by a CAD system, 1600. The CAD system output, 1620, typically includes full or reduced resolution versions of the input imagery with CAD detections overlaid on the images. In one embodiment, the technologist obtains a printed version of the CAD output. In this case, the CAD output paper is stored with the films, and forwarded to a radiologist for use in reading the films for signs of disease at step 1640. In another embodiment, the CAD system output is shown on an electronic display, including a monitor, a wireless hand held tablet, or a personal digital assistant. In this case, the radiologist retrieves the CAD output imagery associated with the films being reviewed by providing identifying information to the Image and CAD output database.

After the initial reading, the radiologist refers to the paper or electronic CAD output, which shows CAD detections overlaid on the medical images. Further action is required when the radiologist detects a suspicious region. Typically the further action involves the radiologist requesting additional imaging from the technologist to provide more information about a specific location in the body. The additional imagery does not necessarily have to be x-

ray. Other modalities are ultrasound, magnetic resonance imaging, and computed tomography. The radiologist annotates the CAD system output in step 1660, producing the annotated CAD output, 1680. Importantly, the radiologist may mark any area of the output imagery and is not restricted to marking only those areas detected by the CAD system. The technologist then uses the annotated CAD output in a subsequent procedure, 1699.

The method of the invention supports non-networked and networked clinical environments. In a non-networked environment, where the CAD outputs are printed on paper, the radiologist identifies the specific area by pencil or pen marks in the context of anatomical detail. The annotated paper CAD output is then transferred to the technologist, who may be remotely located from the radiologist. In another embodiment for a networked environment, the annotation information is transferred through the network and stored in the database. The technologist any authorized network user can access the imagery, CAD system outputs and radiologists annotations.

When the CAD system output is printed on paper, the radiologist uses a marker to identify the area on the printout. Fig. 36 shows a paper-based CAD system output. The radiologist marks of Fig. 37 are in addition to the CAD detections and drawn directly on the medical images of the CAD output. The annotated paper output may then stored with the films in patient's folder. In the case of electronic CAD output, the radiologist workstation accepts input information from the radiologist. In a preferred embodiment, the medical images are displayed on a touch screen monitor, Fig. 38. The radiologist touches an area on the image to highlight a suspicious location as shown in Fig. 39. In Fig. 39, the indications made by the radiologist are shown with the + symbol. Each of the CAD detections is shown as marked by the radiologist in Fig. 39. The annotation information is stored in a database. In an alternative embodiment, the radiologist uses a computer mouse to identify image areas.

In addition to location information, a radiologist may provide other descriptive characteristics regarding the suspicious area. Associated information regarding the suspicious area is preferably provided through a graphical user interface, including pull-down menus, as shown in Fig. 40 and Fig. 41. Commonly included information comprises type of lesion, BIRADS rating, and size. Additionally, an annotation stating that the radiologist "noted and dismissed" a CAD mark is useful. This annotation provides a definitive record that the CAD output was evaluated and deemed not suspicious enough to warrant further study.

An alternative method for providing annotation information includes recording an audio file of the radiologists spoken comments. The radiologist selects a region to annotate, starts the recording session by touching an “record button”, then provides spoken details regarding the selected region. Details may include diagnostic comments and instructions for next the next procedure to be performed on the selected region. The audio file is stored in association with the CAD information and is available to the technologist.

When the technologist performs the re-imaging procedure, the imagery annotated by the radiologist is retrieved from the database via the network for electronic CAD output, or from the patient folder in case of paper CAD output. The technologist then refers to the annotated imagery to guide the re-imaging procedure. The subsequent imaging procedure typically provides greater resolution images across a smaller physical area. Therefore, accurate localization of the indicated suspicious area is essential for obtaining relevant diagnostic information. The efficiency and diagnostic quality of the subsequent imagery is improved by virtue of the specific location and reference information available from the CAD output. That is, since the annotated CAD output contains a mark from the radiologist at the suspicious area overlaid on the originally imaged body part, all existing anatomical landmarks may be used to help the technologist optimally position the patient. The subsequent imaging procedure can be the next scheduled screening exam, commonly occurring one to two years after the current exam. In this case, the annotated CAD output will show the technologist what was of concern in the prior exam.

For workflow efficiency, it is desirable for the CAD systems to display marks only at locations not previously considered by the radiologist. Therefore, an additional aspect of the invention is in using radiologist supplied detections to modify the displayed CAD system detections. When the radiologist notes a suspicious area, they provide an electronic description including the type of lesion and its location. Prior to displaying the CAD system output, the radiologist detections are compared to the CAD system detections. Only the CAD system detections that do not include the radiologist detections are displayed.

### Radiologist Modification of CAD Detections

In some situations, the radiologist may desire to create a modified version of the

CAD output. For example, the CAD system may detect one suspicious area but additionally produce marks that the radiologist decides are diagnostically irrelevant. In this case, the radiologist may choose to delete the marks deemed irrelevant to prevent possible confusion during the re-imaging procedure. Preferably, the dropped CAD marks are annotated as “noted but judged as of no concern”. In another embodiment, the radiologist may use the CAD output to provide feedback to the patient. Again, consider a patient whose CAD output contains one true mark but several extraneous false marks. The radiologist may provide the patient with the CAD output in explaining the need for an additional imaging procedure. However, the patient may become unduly concerned by the presence of the extraneous false marks. Therefore, it is advantageous for the radiologist to have the capability to modify the CAD output such that only those marks deemed suspicious are visible in the version sent or shown to the patient. In yet another embodiment, consider the situation when a suspicious area is not detected by the CAD system. In this case, the radiologist may add a detection marker to the CAD output imagery.

Referring again to Fig. 34, a modified CAD output is created as follows. The radiologist enters information at the workstation 1800 and retrieves the CAD output from the database, 1400. The CAD imagery and output are displayed at the workstation. Fig. 42 shows a CAD system output with three CAD detections. Fig. 43 shows the modified CAD system output with only two displayed detections. The radiologist selects particular CAD marks to remove by means of a touch screen or mouse. Alternatively, only the marks to retain may be specified.

Marks to add may also be specified. Fig. 44 shows a CAD system output having only 1 CAD detection. Fig. 45 shows a radiologist added mark in the RCC view. In all cases, a modified CAD output file is created and stored in the database. The modified CAD output can then be printed or displayed at the technologist workstation by way of commands from either the radiologist or technologist workstation. The original CAD output remains unmodified.

In facilities equipped with electronic reporting systems, the modified CAD output and annotations made by the radiologist may be incorporated into electronic reporting system records.

What is claimed is: